Computational Autonomous Mental Development:
A White Paper for Suggesting a New Initiative

Juyang Weng[1], James McClelland[2], Alex Pentland[3], Olaf Sporns[4],
Ida Stockman[5], Mriganka Sur[6] and Ether Thelen[7]

Abstract

A new synthesis of the neural, behavioral, and computer sciences is on the horizon. The topic that promises to unite these disparate fields is computational autonomous mental development. The term "mental" refers to cognitive, behavioral and other mental skills that are exhibited by humans, higher animals and artificial systems. Computational autonomous mental development refers to the computational process by which a brain-like machine, natural or artificial, develops mental skills under the guidance of an intrinsic *developmental program* and through its own *autonomous* activities using its sensors and effectors to interact with its environment. The developmental program for an animal resides in the genes as a result of many generations of evolution; while that for a machine is initially programmed into the machine by humans but the environment changes the ways that the developmental program operates. The synthesis is inspired by new discoveries in neuroscience that highlight the exquisite plasticity of the brain with experience through infancy and adulthood, by new theories and computational modeling of human cognitive development, and by methodological and computational advances in AI and robotics that make it possible for machines to autonomously develop their own intelligence. Potentially, there are enormous benefits as a result of this synthesis: For behavioral and neural scientists, it promises a deeper, more precise and more systematic understanding about the ways our brain works through the computational study of its developmental processes. For the engineering and computer sciences, there is the vision of greatly enhanced capability for machines to interact with humans and to process information to a degree that requires kinds of machine intelligence other than those possible before.

[1] Department of Computer Science and Engineering, Michigan State Univeisity, East Lansing, MI 48824. Tel: 517-353-4388, Mail: weng@cse.msu.edu.

[2] Center for the Neural Basis of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213. Tel: 412-268-3157, Email: jlm@cnbc.cmu.edu.

[3] The Media Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139-4307. Tel: 617-253-0648, Email: sandy@media.mit.edu.

[4] Department of Psychology, Indiana University, Bloomington, IN 47405. Tel: 812- 855-2772, Email: osporns@indiana.edu.

[5] Department of Audiology and Speech Sciences, Michigan State University, East Lansing, MI 48824. Tel: 517-353-6764, Email: stockma1@msu.edu.

[6] Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139-4307. Tel: 617- 253-8784, Email: msur@wccf.mit.edu.

[7] Department of Psychology, Indiana University, Bloomington, IN 47405.Tel: 812-855-2042, Email: thelene@indiana.edu.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **03 OCT 2000**	2. REPORT TYPE	3. DATES COVERED **00-10-2000 to 00-10-2000**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Computational Autonomous Mental Development: A White Paper for Suggesting a New Initiative	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Michigan State University,Department of Computer Science and Engineering,East Lansing,MI,48824**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**16**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

The purpose of this white paper is to provide essential but concise information to the most relevant branches of the US government about this new direction so that timely actions can be taken to take advantage of the new development. This white paper contains the following information:

1. A brief executive summary of the background of this subject.
2. Major characteristics of the new direction.
3. The predicted impact.
4. What the government can do now.

1 Background

This white paper is a direct outcome of the Workshop on Development and Learning, funded by NSF and DARPA, held at Michigan State University, East Lansing, MI, April 5 - 7, 2000 (http://www.cse.msu.edu/dl/). This workshop was attended by about 30 distinguished researchers in neuroscience, developmental psychology, machine intelligence and robotics who are working on related subjects in their fields. This workshop was organized at a time when an increasing number of researchers have realized the importance of computational study of cognitive development and the need for a multidisciplinary forum. The goal of this workshop was to discuss the state-of-the-art in research on cognitive development and to discuss, initiate and plan future research on this subject. After extensive discussions at the workshop, the participants reached a consensus that a white paper for related US funding agencies and the related research communities is needed.

In the last 50 years, the scientific communities have made significant progress in understanding the ways human mind works, the power and limitation of existing machines, as well as the relationship between humans and machines. It is now clear that a *developed* human mind, that of a normal human adult, is extremely complex. It is also clear that the early optimism in the 60's and the 70's about quick progress in artificial intelligence (AI) such as vision, speech, and language, was not well founded, at least not so within the realm of traditional approaches that have been extensively experimented with so far. However, past work using traditional approaches is by no means unimportant. In fact, they are the incubator for the birth and growth of the new direction for machine intelligence: the direction of autonomous cognitive development. As S. Kuhn wrote in his book titled ``The Structure of Scientific Revolution" [1]: "Because it demands large-scale paradigm destruction and major shifts in the problems and techniques of normal science, the emergence of new theories is generally preceded by a period of pronounced professional insecurity. As one might expect, that insecurity is generated by the persistent failure of the puzzles of normal science to come out as they should. Failure of existing rules is the prelude to a search for new ones." (page 68).

The puzzle pieces from recent advances in related fields are beginning to reveal a picture of *cognitive development,* which is rooted in a deeper understanding of the computational mechanisms that give rise to mind as opposed to the working of a static mind itself. We briefly summarize below some of these new thought-provoking advances.

1.1 Neuroscience and psychology

A traditional view about the human brain is that its function is very much pre-determined by the genes. In this view, the brain unfolds its pre-determined representation during development,

which starts from the time of conception. Each area of the brain has fully pre-determined representation that is specific to the functions of the area. However, recent advances in brain plasticity have begun to reveal a very different picture of brain development.   For example, researchers at MIT [2] have discovered that if the optic nerves originating from the eyes is rewired into the auditory cortex of an animal (a ferret) early in life, the auditory cortex gradually takes on a representation that is normally observed in visual cortex.   Further, the animals successfully learned to perform vision tasks using the auditory cortex.   In other words, the rewired ferrets can see using the brain zone that is normally assigned for sound. This discovery seems to suggest that the cortex is governed by self-organizing mechanisms, which derive representation and architecture according to the input signals, either visual or auditory. As another example, studies by researchers at the University of California at San Francisco [3] have shown that the finger skin areas from which a neuron in somatic cortex receives sensory signals (called receptive field of the neuron) can change according to sensory experience.   If multiple fingers of the adult monkeys receive consistent synchronized pulse stimuli from a cross-finger bar for several days, the receptive field changes drastically, from covering only a single finger in normal cases to covering multiple fingers.   This result appears to indicate that the self-organizing program of our brain automatically selects the source of sensory input within a candidate area according to the statistical properties of the actual sensory signal that is received.   These and other related studies on brain plasticity prompt us to rethink the traditional rigid view about the brain.   It appears that the developmental program in the genes does not rigidly determine brain architecture and representation.   Of at least equal importance are epigenetic factors, such as statistical properties of the sensory signals that are received and how these signals are used to derive the representation and architecture of the brain.

In recent years, computational modeling of neural development is becoming an active subject of study in neural science and psychology. For example, there are several computational models for the development of response patterns in the retina, the lateral geniculate nucleus, and simple cells in the visual cortex. A particular subject that is now actively investigated concerns the mechanisms for developing orientational selectivity in the simple cells of the visual cortex. Some recent work in psychology has started to explain the global process of cognitive development using architecture of networks [4]. Another new trend in psychology is to use explicit dynamic models to explain some well-known facts about infant behaviors (e.g., the work at Indiana University [5]). These quantitative studies have begun to produce results that are more explicit and empirically verifiable than vague verbal theories and arguments. Psychology has begun to move from qualitative descriptive models to more rigorous quantitative models for studying cognitive and behavioral processes. However, the mainstream in psychology is still to explain phenomena about ways the brain works but not *computational mechanisms* of their development. The area of developmental psychology has produced many important results, but it needs *computational* studies about mental development.   As discussed above, the fields of neuroscience and psychology have recently produced some computational models of developmental mechanisms, but they are about early processing (early in the order of processing steps in the brain, such as orientation sensitive cells).   The new research direction is to work out *experimentally implementable computational models* about mental development that can explain not only earlier processing, but also later processing in the brain, while these models are increasingly simulated and tested on robots.

1.2 Robotics and Machine Intelligence

Although cognitive development in humans is a well-known fact, the counterpart for machines did not receive serious attention until the mid 90's. It has long been believed that the approach to machine intelligence does not have to follow what human minds do, just as modern airplanes do

not fly like birds.   Gradually, many AI researchers started to realize that machine intelligence requires much more cognitive and behavioral capabilities than most had realized.   Flying is a very simple problem in comparison with machine intelligence.   Furthermore, many AI researchers have already realized that machine intelligence requires "grounding" --- concepts must be grounded in real sensory experience about the physical world, which in turn requires the machine to have a sensor-rich body (i.e., embodiment) that can directly sense the physical world and act upon it.   Given that rounded sensing and action, including learning, have been extensively studied in robotics for many years, why then does the reality of highly intelligent machines still seem so remote?   Recently, it was pointed out [6] that what has been sorely missing from machines is *autonomous mental development*.   The following diagram relates mental development by machine with that of a human.
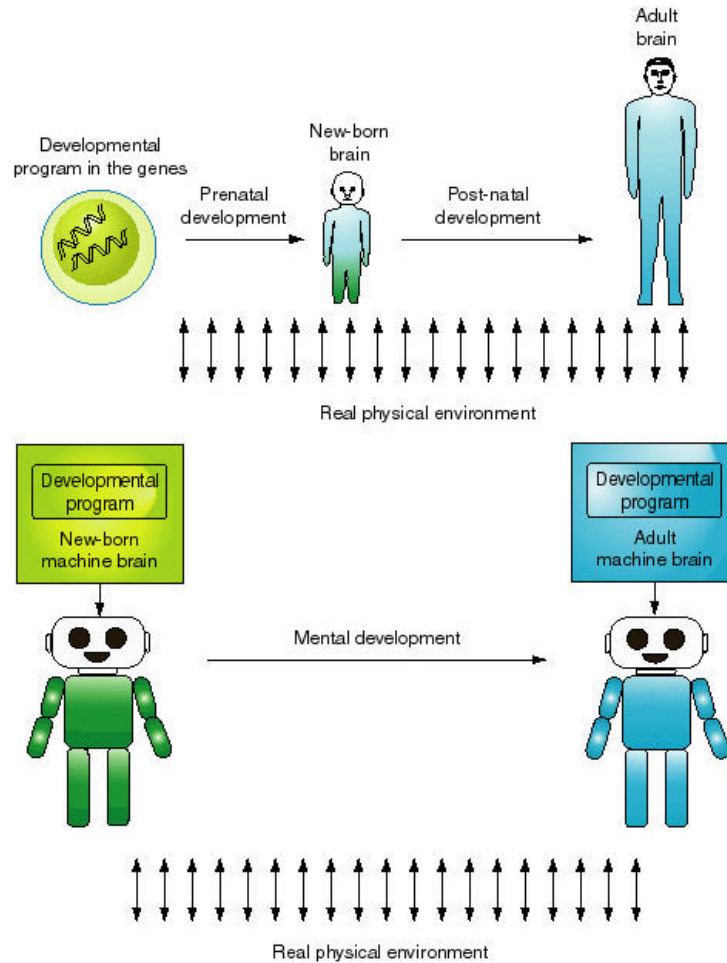

Figure 1: The major issue of computational study of autonomous mental development, both in humans and machines is to study the computational principles in the developmental program and the process of autonomous mental development through online, real time interactions with the physical environment.

Autonomous mental development requires a fundamental change in the basic way engineering has been done (i.e., paradigm) traditionally.   The current *manual development paradigm* is as follows:

1. *Start with a task:* Given a task to be executed by a machine, it is the human engineer who understands the task (not the machine).
2. *Design a task-specific representation:* The human engineer translates his understanding into a representation (e.g., giving some symbols or rules that represent particular concepts for the task and the correspondence between the symbols and physical concepts). The representation reflects how the human engineer understands the task.
3. *Programming for the specific task:* The human engineer then writes a program (or designs a mechanism) that controls the machine to perform the task using the representation.
4. *Run the program on the machine.* If machine learning is used, sensory data are then used to modify the parameters of the task-specific representation. However, since the representation is designed for the specific task only, the machine cannot do anything beyond the pre-designed representation. In fact, it does not even know what it is doing. All it does is run the program.

The new paradigm, *the autonomous development paradigm,* for constructing developmental machines or robots, is as follows:
1. *Design body:* According to the general ecological condition in which the robot will work (e.g., on-land or underwater), human designers determine the sensors, the effectors and the computational resources that the robot needs and then design a sensor-rich robot body.
2. *Design developmental program:* A human programmer designs a *developmental program* for the robot.
3. *Birth:* A human operator turns on the robot whose computer then runs the developmental program.
4. *Develop mind:* Humans mentally "raise" the developmental robot by interacting with it. The robot develops its mental skills through real-time, online interactions with the environment which includes humans (e.g., let them attend special lessons). Human trainers teach robots through verbal, gestural or written commands in much the same way as parents teach their children. New skills and concepts are autonomously learned by the robots everyday. The software (brain) can be downloaded from robots of different mental ages to be run by millions of other computers, e.g., desktop computers.

A system for mental development is not simply an incremental learning system that can grow from small to big in terms of its memory size. Such systems have already existed (e.g., some systems that use neural network techniques). Traditional machine learning systems still operate in the *manual* developmental mode outlined above but mental development requires the *autonomous* developmental mode. What is the basic difference? In a traditional machine learning mod, the task to be executed by the machine is known that thus the programmer can design the representation for the task in his programming. But in the new developmental approach, the programmer of the developmental approach do not know what tasks that the robot will execute and thus the representation and typically also architecture of the representation have to be autonomous generated online in the later learning stage. In other words, programming for a developmental program is not task-specific. A developmental program must be able to generate *representation and architecture* autonomously online directly from raw sensory signals.

Computational study of autonomous development of representation and architecture is a relatively new area of study, for both the human brain and robot brain. With this capability of generating new representation, machines are able to learn subjects that their programmers do not know, or have not even thought about, just as human children learn subjects that their parents do not know about. The "subjects" here do not have to be all at high abstraction levels. For example, a developmental program must autonomously learn to "see" things that the programmer has never seen or has even thought about during programming. This task-nonspecific nature of

developmental program has far-reaching implications for developing complex mental capabilities. For instance, researchers in the computer vision field have been facing great challenges in their attempts to write vision programs that can deal with unknown objects in unknown environments.

The essence of autonomous mental development by machines is the capability of learning directly, interactively, and incrementally from the environment using its own sensors and effectors. Therefore, a computer with only impoverished sensors and effectors cannot do mental development well. A neural network that can only accept human edited offline sensory data does not develop its mind either, even if it can learn incrementally. A *developmental robot* is a robot that runs a developmental program and is allowed to learn and practice autonomously in the real physical world. The following figure summarizes the basic differences between the traditional engineering paradigm and the new developmental paradigm. The characterization in the table for a developmental program is also true for the natural developmental program of higher animals. Although human infants do exhibit some innate behaviors at the birth time (i.e., sensor- and effector specific), what actual tasks that a human individual will learn after the birth is unknown at the conception time. For example, what toys that the infant will play is not known at that time.

Properties of Program	Traditional Programs[8]	Developmental Programs
Sensor-specific and effector-specific	Yes	Yes
Program is task-nonspecific	No	Yes
Tasks are unknown at programming time	No	Yes
Generate representation automatically[9]	No	Yes
Animal-like online learning[10]	No	Yes
Open-ended learning of more new tasks	No	Yes

Although the concept ``developmental program'' for machines is relatively new, many techniques useful for developmental programs have already been created in the past several decades in the fields of pattern recognition, robotics and machine intelligence. Although these techniques were not designed originally for automatic mental development, they are very useful for developmental programs if they are applicable to high-dimensional data. These techniques are being used in very innovative ways in designing developmental programs. Several developmental programs have been designed and tested on robots. Running a developmental program, the robots interact with the environment in real time using their sensors and effectors. Internal representation, perceptual capabilities and behavioral capabilities are developed autonomously as a result of interaction between the developmental program and the environment. Humans, as a part of the environment, interact with such robots only through the robot's sensors and effectors. Just like the nature-nurture interplay for human mental development, the cognitive and behavioral skills of such a robot result from extensive interaction between what is programmed ("innate" developmental program) and what is sensed through real-time online experience. Mind and intelligence emerge gradually from such interactions.

Early examples of such developmental robots include Darwin V robot at The Neurosciences Institute, San Diego and SAIL robot at Michigan State University, developed independently

---

[8] Including traditional learning programs such as traditional artificial neural networks.
[9] For tasks unknown at the programming time.
[10] Can learn from one instance in real time; learn while performing.

around the same time but with very different goals. The goal of Darwin V [7] was to provide a concrete example for how the properties of more complex and realistic neural circuits are determined by the behavioral and environmental interactions of an autonomous device. Darwin V has been tested for the development of generalization behaviors in response to visual stimuli at different positions and orientations (visual invariance learning). It also has been tested for the association of aversive and appetitive stimuli with visual stimuli (value learning). SAIL was designed as an engineering testbed for developmental programs that are meant for scaling up to complex cognitive and behavioral capabilities [8]. The SAIL-2 developmental program has been tested for autonomous derivation of architecture and representation through online, real time development of association (1) between visual stimuli of objects and eye aiming for the objects (object evoked visual attention); (2) between visual stimuli of objects and arm pre-reaching for the object (vision evoked object reaching); (3) between voice stimuli and arm actions (verbal command learning and execution) and (4) between visual stimuli and locomotion effectors (vision-guided navigation). Other studies of online learning are directed towards fully autonomous developmental systems. For example, the work at MIT has associated video images of an object with a synchronized voice (pronounced verbal name of the object) [9]. The work at the University of Massachusetts at Amherst has investigated the use of coupling among robot leg joints that have been observed in human infants to reduce the search space for a desirable turning gait [10]. Although the history of developmental robotics is very short, some experiments examplified by the above studies have demonstrated capabilities that have not been achieved before or are very difficult to achieve using traditional methods, such as visual object recognition, verbal communication, hand-eye coordination, autonomous navigation, value acquisition (i.e., learning the value of actions), and multimodal association in real time. We are aware that more groups in the US and other countries have started investigations that reflect or are consistent with this new direction.

2. Some Major Characteristics of Research on Mental Development

2.1 More tractable

It is known that a developed adult human brain is extremely complex, as an epigenetic product of long-term and extensive interactions with the complex human world. That is a major reason why all the major fields that are related to intelligence have met tremendous challenges, including neuroscience, cognitive science, artificial intelligence (especially vision, speech, languages, planning, decision making, etc), robotics and philosophy, on related subjects. However, the developmental principles for the brain in the complex human world, however, should not be as complex as the human world itself. For example, the visual world is very complex, but the developmental principles that are used by the brain to derive various filters for processing visual signals should not be as complex as the visual world itself. Due to the task nonspecific nature of a developmental program, designing a developmental program needs only to model developmental mechanisms for signals, relieving human from programming for the complex world that the signal represents. Therefore, the computational study of mental development could be more tractable than traditional approaches to understanding natural intelligence and constructing intelligent machines.

2.2 Unified framework

Studies of mental development will establish a unified framework for our understanding of a wide variety of cognitive and behavioral capabilities. Discovery of mechanisms responsible for

developing cognitive and behavioral capabilities in general settings requires more systematic work than an account of a particular individual capability in a special setting (e.g., visual attention selection in a particular room). The sharing of common developmental principles by visual and auditory sensing modalities, as recent neuroscience studies have demonstrated, will encourage scientists to further discover underlying developmental principles that are shared not only by different sensing and effector modalities, but also by different higher brain functions.

Traditionally, it was considered that processing methods for vision and speech should be very different, in both humans and machines. For the same reason, traditional methods for different AI problems are typically very different, resulting in what is well known now as the fragmentation of the AI field. Potentially, AI can be applied to all possible areas of human life and each application area can lead to a fragment of AI if it is treated in an *ad hoc* way. The unified computational framework of mental development will fundamentally change the current fragmented landscape of AI in the years to come, since different applications are expected to correspond to different lessons that can be taught to the same developmental robot at different mental ages. We should also see much more interactions and collaborations among scientists and engineers in neuroscience, psychology, robotics, AI and other related fields, due to the very similar research issues these fields face under the theme of autonomous mental development.

2.3 Task-nonspecific

In contrast to the task-specific nature of the traditional engineering paradigm in AI, developmental programs for machines will be *task-nonspecific*. The power of a developmental program is its general applicability to many different tasks. A developmental program may contain certain pre-processing stages that are specific to some type of sensors or effectors, such as camera or touch sensors. In this sense, it is body-specific or species-specific. However, it is not task-specific. A developmental program can be run to develop skills for many different tasks, with simpler skills being learned to prepare skills for learning more complex skills. Recently, the scientific community has gained a more complete understanding of human intelligence. As Howard Gardner put it in his book titled *Multiple Intelligences* [11], human intelligence has multiple facets, including linguistic, logical-mathematical, musical, bodily-kinesthetic, spatial, interpersonal, and intrapersonal intelligences. This view points out a very rich ensemble of inter-related cognitive and behavioral capabilities that give rise to human intelligence. The same is true for machine intelligence. Any particular capability that we regard as intelligence in a general setting, such as the visual capability of recognizing various persons on a busy street or the language capability of talking about technology, is not an isolated single entity. It requires the support of many skills developed through extensive real-world experience via sensors and effectors.

2.4 Computational

Further, developmental mechanisms seem to be very much quantitative in nature and thus require clear *computational models*. We will see more complete computational models for mental development that can be simulated by computers and robots for many different environmental conditions and the results can be verified against studies about humans. We will see more efforts on computational modeling of mental development, for humans and machines, that are clearly understandable, implementable on machines and can be subject to rigorous verification and comparison. This will reflect the maturation of the related fields.

2.5 Recursive and active

Development discourages any static or rigid view of the mind. A developed human mind is a snapshot of many years of *recursive and active* mental construction by the developmental program in the human genes, utilizing the sensory and action experience through life time. The term "*recursive*" means that later mental development relies on the cognitive and behavioral capabilities that have been developed earlier. The term "active" means that each individual is actively involved in constructing knowledge online by changing and perceiving the environment via interaction. Different actions lead to different experiences. The same is true for developmental robots. The recursive and active nature of development discourages the approach of collecting offline data and spoon-feeding them into a machine, which is a prevailing practice in current machine learning studies. Sensory data cannot be pre-specified since what sensory data are sensed depends on the online actions executed in real time and the history of experience. For these reasons, a sensor-rich and effector-rich mobile robot seems more suited as a testbed for mental development than a sensor-impoverished and effector-impoverished desk-top computer.

2.6 Developmental capabilities as unified metrics for machine intelligence

The *criteria for measuring machine intelligence* will fundamentally change. The metrics that can be used to measure the power of such a new kind of machine is primarily their autonomous interactive learning capabilities in complex human environments. In other words, it is the capability of mental development instead of what the machine can do under a pre-specified setting. Such performance metrics can be adapted from those used by clinical psychologists for testing the mental development of human infants (e.g., The Bayley Scales of Infant Development [12]) and children (e.g., The Leiter International Performance Scale [13]). The mental age that is used for measuring human intelligence in these tests can be adapted to a scale for measuring machine intelligence. This is a fundamental improvement over the current metrics that measure what a machine can do under a specific setting and task condition. What a machine can do in a specific situation is the intelligence of the machine programmer, not the machine itself. For example, an interactive dictionary stores a lot of human knowledge and it can do remarkable things for humans, but it is not intelligent. Test criteria for machine intelligence may also provide quantitative feedback for improving the intelligence tests for humans.

3 Predicted Impacts

The history of science and technology has shown that impressive technical improvement and persistent cost reduction will follow an important scientific revolution. The amount of technical improvement and cost reduction can be so great that it was difficult to foresee at the time of actual revolution. Two well known examples are the internal combustion engine technology leading to today's automobiles and the Von Neumann machine idea and the semiconductor technology leading to today's computers. The following predictions may seem to be overly optimistic today, but time could prove them to be true.

3.1 Understanding of human mind

The impact on the scientific understanding of human mind will be far reaching. This new endeavor will potentially improve our understanding about one of the most complex subjects that face mankind today --- our own minds. For example, what are the basic mechanisms that

govern the ways in which the mind develops?   To what degree can the environment change the formation of the mind?   What can the environment do to effectively and positively influence the human mind and improve the life of mankind?   The study of mental development is to search for the root of the mind. Without studying the computational models of mental development, these questions cannot be sufficiently and clearly answered.

3.2 Human life

Success in creating machines capable of autonomous mental development will likely improve the quality of human life in several different ways. Humans have not been very successful yet in programming mental capabilities into a machine, such as vision, speech and language. Developmental robots could show their strength in developing these capabilities. When these capabilities are developed to near the level of higher animals, developmental robots can be used as human assistants, from factories to households.   Computers would have a much improved interface with humans.   Their capability of communicating with humans using visual, auditory and touch cues will improve greatly since these are common sensory modalities of developmental robots.   Developmental robots will attend special robot classes to be trained for different, simple, repetitive cognition tasks. For ordinary daily life, developmental robots will learn to perform tasks that human do not like to do all the time, such as screening emails, browsing documents, mowing a lawn or watching a baby.   For works in demanding environments, such as undersea searching, space exploration, cleaning up of nuclear waste, microscopic manipulation, border surveillance, the developmental robots will extend the physical limitation of humans. In entertainment and education, developmental robots can serve their personality roles in mass media, museums, and amusement parks.

Why did all these developments not occur in the past or do so very successfully? Traditional AI did not pay sufficient attention to, or was not serious about, autonomous mental development for machines until just a few years ago. Currently, all the efforts for building AI systems follow the manual development paradigm, with a few recent exceptions mentioned above.   They have produced special purpose machines, instead of autonomous general-purpose learners. With the new autonomous development paradigm, human programmers are not required to write a particular program for each of the tasks that they want the machines to perform, a task that is extremely difficult if it requires what we consider as intelligence. Instead, what the human programmers need to do is to write a general purpose developmental program. Although developmental programs are by no means easy to design, they seem easier for humans to understand and to improve than many special systems designed for specific AI tasks.

The practical use of developmental robots also rests on the ease of training. The user of a developmental robot does not need to write a program or manually feed data when he wants to teach the robot. He trains the robot the way a human child is taught --- by showing what to do and how to do while talking to it, encouraging or disapproving from time to time.   Thus, anybody can train a highly improved developmental robot of the future, a child, an elderly person, a teacher, a worker and so on. This is the basic reason why developmental robots could be popular. Computers would not be so popular today if they were not as easy to use as today's computers which have very intuitive graphic user interface.

3.2 Economy and jobs

The economic impact of developmental robots will depend on the market size of developmental robots. The country that takes the lead in developmental robots will first create a new industry

for this new kind of machines. This new industry will take advantage of the advanced automobile industry to develop sensor-rich humanoid robots (Honda in Japan has already started it).   The computer industry will take advantage of the need for developmental robots to build computers and memories best suited for the computational need of developmental robots.   The cost for large storage should drop consistently when the computer market continues to grow. For example, the cost of hard-disk storage that is of human brain size in terms of number of bytes has already dropped to around $250,000 today (June 2000).   Real-time speed with large memory is reached through coarse-to-fine memory search schemes.   There will be a new industry for humanoid robots, fueled by the need for building bodies for developmental robots.   Many different types of bodies, designed for different working conditions and environments will be made to satisfy the increased application scope of developmental robots.   It is expected that in the next 10 to 20 years, the developmental robot industry will primarily aim at professional applications, for research institutions, amusement parks, public service areas, and the defense industry.   During this period, consumers can benefit from using the software that is downloaded from professional robots. Eventually, developmental humanoid robots are expected to cost the same as a car plus a high-end personal computer. In that time, this new industry could be as large as the automobile and computer industry combined today. The country that takes the lead in this new endeavor will create an abundance of economical activities and well-paid jobs related to this new industry.

3.4 Medicine

The knowledge created by this new endeavor will also improve medical care.   It will provide basic knowledge useful for treating learning disabilities, mental disorders, and mental problems associated with aging.   For example, what developmental mechanisms are responsible for attention deficiency?   What developmental mechanisms are responsible for an individual to establish the value of an event, a behavior, or the social norm?   What techniques are effective for teachers to improve the development of certain cognitive and behavioral capabilities? Computationally, which areas of the brain are responsible for certain mental disorders? During aging, which brain mechanism is likely to deteriorate first and what remedies are possible?

4 Suggestions to US Government Funding Agencies

It takes vision and courage to embark on a new endeavor. The Internet grew out of ARPANET whose seed was planted in 1969 by a DARPA supported project titled Resource Sharing Computer Networks. Funding from the US Federal Government played an important role in the development of ARPANET and Internet. By the time the last piece of NSF backbone ceased to work in the Internet on April 30, 1995, the Internet had become an outstanding example of how the government can play a positive role in facilitating the growth of a new industry. Thanks to the popularity of the Internet, the number of computers sold annually in this country has exceeded the number of automobiles sold, and the Internet has entered the homes of millions of American families.   Its world-wide use coverage is growing very fast as well.

4.1 A well-focused new program is needed

Timely Federal funding support is crucial for this new direction at its fledging stage, especially at a time when the old paradigm of doing engineering is still deeply entrenched in the research community.   A new well-focused program is needed to guide research in this area --- a program

similar to the NSF program on nanotechnology.   Since this is a multidisciplinary effort that has a high impact on science, engineering, health and defense, it appears that NIH, NSF, DARPA and other funding agencies could all be involved.   A suggested title for this new program is "Computational Autonomous Mental Development." We propose the following central themes for the new Autonomous Mental Development Program:

1. For fields that have humans as research subjects, i.e., neuroscience and psychology, the goal is to study the computational principles underlying the development of the human brain. The goal is not to study just phenomena or mechanisms of a developed brain, since this kind of effort can be supported by other existing programs.   The goal is not to study noncomputational aspects of development either, since there are programs that support such studies in neuroscience and developmental psychology.
2. For fields that have machines as research subjects, i.e., robotics and machine intelligence, the goal is to realize autonomous mental development by robots for performing various tasks that programmers do not know at the time of programming.   The program is not designed to look for methods with a given known task, since there are existing programs that support studies on such subjects.   The goal is not to study methods that does not enable robots to learn directly through sensors that sense the real physical world and through effectors that act on the real physical world.   These methods do not fully support autonomous mental development and can be supported by existing programs.
3. The program encourages collaboration among fields that study human and machine cognitive and behavioral development. Biologically motivated mental development methods for robots and verification of biological models of mental development for robots are especially encouraged.

Since this is a new endeavor, it is expected that some members of the proposal review panel have not yet gained full appreciation of this new program.   It is possible that a consensus is more easily reached for supporting a middle ground work than a truly revolutionary work. Therefore, it is very important for the program announcement to clearly state the goals of the new program, especially emphasizing the basic difference between this new direction and the traditional ones --- the developmental program is in place before the tasks are known, as outlined in the above table. Effective measures should be put in place to ensure that the reviewers evaluate the proposals strictly based on the goals of the program, which may seem challenging or risky according to their own conventional wisdom.

According to the funding levels of similar programs in the past, we suggest that the program support two research centers at about $3M per year for each center for 5 years. Each of these two centers should involve Co-PIs from the human intelligence side and from the machine intelligence side.   The center type efforts are needed because this research endeavor requires systematic testing of diverse types of cognitive and behavioral capabilities and thus it requires coordinated efforts involving a significant number of Co-PIs.   A center can involve several institutions. Another $6M per year for 5 years is for individual projects.   These individual projects can focus on specific aspects of autonomous mental development. $1M per year for 5 years, $10M in total, is needed for each of the two industrial partner that works closely with the two centers to develop commercial sensor-rich and effector-rich humanoid robots suited for autonomous mental development.   Humanoid robots for mental development require rich sensors and effectors and they also must be mobile. These requirements have not been adequately met by existing robots, conventional or humanoid. This amounts to $70M in total for a 5-year program on mental development.

4.2 Why now?

As we discussed above, recent discoveries about the human brain tell us loud and clear that our brain utilizes developmental principles that are shared by different sensing and effector modalities.   Since higher brain functions appear to be even more plastic than early sensory processing functions, it is expected that higher brain functions also use developmental principles that are generally applicable to different subject matters that humans learn.   The time is right to study what these developmental principles really may be.

Technically, it is now possible to study massively parallel, distributed brain activities and relate them to mental development. The advances in brain imaging techniques, such as fMRI, MR, PET, and ERP now allow high resolution, concurrent, and real-time measurement of brain activities.

In the machine intelligence and robotics fields, the fundamental difference between the way human mind develops and the traditional engineering paradigm for machine development was recently identified as a fundamental reason for the difficulties in AI. Studies about the fundamental limitations of the current engineering paradigm have recently started.   Some preliminary computational models for developing the mind by machines were recently proposed and tested.   These early efforts have achieved some encouraging results that have not been possible before using the traditional engineering paradigm. Therefore, computational models of mental development for machines are not beyond human comprehension.   They are within the manageable scope for humans to model computationally.

The performance-to-cost ratio of computers has reached a level that makes it practical to simulate brain development in real time on a robot. For example, the development of the most computational challenging modality, vision, can now be simulated on real robot in real time using software run on a PC workstation.

Technology for building robots has also been improved significantly.   In recent years, research laboratories and related industries in US and Japan gained remarkable experience in actually building robots that resemble human and animal bodies with similar articulate structures, from human-size humanoid robots (e.g., the series of Honda humanoid robots) to advanced consumer toy robots (e.g., Sony AIBO dog robots).   The robotic technology is ready for building various humanoid or animal robots as bodies for developmental machines.

4.3 Research issues

The task-nonspecific nature of mental development should make the studies of mental development easier than the traditional task-specific approaches.   This is true both for human subjects (neuroscience and psychology) and machine subjects (AI and robotics).   From the computational view of mental development, the research issues center around sensory signals and effector signals with internal states that are automatically generated. A developmental program will associate signals that are from different sensors, stored in internal status and sent to effectors, but its programmer does not need to know what those signals actually mean. To put it intuitively, it is easier to model how an interactive program looks up words from its word memory than to model how the meanings of words in The Merriam-Webster's Dictionary relate to one another. The former is like what a developmental program does for many tasks that a developmental being could come across and the latter is like what all the traditional programs do for a particular task.

To understand this fact better, we take a complex behavior as an example. To model attention selection in a traditional task-specific way, one must understand the nature of the task (e.g., driving a car) and then study the rules of attention selection based on the steps of the task. Such rules are extremely complex (e.g., due to the complex street situation during driving) and the results are ad hoc in the sense that they are not directly applicable to other tasks or even to the same task under different scenarios. In contrast, attention selection by a developmental being is just a part of behaviors that are being developed continuously and constantly. As long as the effectors for attention selection are defined for the body (external effectors) and the brain (internal effectors), the attention selection principles are developed autonomously by the same developmental program in a way very similar to the behaviors from other effectors, such as arms and legs.

Consequently, a series of very interesting and yet manageable new research problems are opened up for fields that study either human or machine subjects. Some of the tractable research problems that can be studied in this 5-year program are suggested below.

1. Schemes for automatic derivation of representations from sensory signals that are sensed from the environment and the body.
2. Schemes for automatic derivation of representation from effector signals, derived from practice experience.
3. Automatic derivation of receptive fields, in both the classic and nonclassic sense. That is, how later processing elements in the brain group outputs from earlier processing elements or sensory elements.
4. Long term memory growth, self-organization and retrieval for high-dimensional neural signal vectors.
5. Working memory formation and self-organization for high-dimensional neural signal vectors. The working memory may include short term sensory memory and the system states.
6. Developmental mechanisms for mediation of conscious and unconscious behaviors. That is, those for mediation among higher and lower level behaviors, such as learned behaviors, learned emotional behaviors, innate emotional behaviors and reflexes.
7. Mechanisms for developing internal behaviors --- those that operate on internal nervous components, including attention selection. This subject includes both developmental mechanisms and training strategies for humans and robots.
8. Attention-directed time warping from continuous states. The time warping issue concerns the time inconsistency between different instances of experience, with the goal of both generalization and discrimination.
9. Autonomous action imitation and self-improvement. The developmental mechanisms underlying an improved behavior pattern that results from individual online instances of related experience.
10. Mechanisms for *communicative learning* and autonomous thinking. Communicative learning refers to learning directly through languages (auditory, visual, tactile, written etc) as children do when they attend classes. These mechanisms are closely related to the development of thinking behavior, which is responsible for planning, decision making and problem solving.

4.4 Expected pay-off of this new program

It is expected that through this 5-year program a series of milestones can be achieved.

1. Global computational models for human mind development will be proposed and tested. These models are not necessarily correct or complete in explaining all the facts about human

mental development. However, they will mark the milestone where humans start to have a clearly explained global computational model of their mental development.

2. Developmental mobile robots will be continuously operational for a period of a half year or longer. This indicates that the developmental programs are stable enough and produce good performance. The performance record at the end of each month after birth will be available for evaluation of the developmental progress.

3. Developmental robots will reach the mental age of about 6 months, as measured by the Bayley Scales of Infant Development. The implication of this milestone is more important than what the mental age implies. Because of the autonomous learning mode of developmental algorithms, such robots will further develop without a change of their developmental programs to acquire more complex behaviors that correspond to higher mental ages. Of course, further improvement of the developmental program is always needed.

4. The first commercial company for developmental robots will have the first product ready for the market, initially for research institutions, amusement parks, industrial applications, and defense.

References

[1] T. S. Kuhn. *The Structure of Scientific Revolution*, University of Chicago Press, third addition, page 68, 1996.

[2] L. von Melchner, S. L. Pallas and M. Sur. Visual behavior mediated by retinal projections directed to the auditory pathway. *Nature*, vol. 404, April 20, pages 871-876, 2000.

[3] X. Wang, M. M. Merzenich and K. Sameshima and W. M. Jenkins. Remodeling of hand representation in adult cortex determined by timing of tactile stimulation. *Nature,* vol. 378, no. 2, pages 13-14, 1995.

[4] J.L. Elman, E. A. Bates, M. H. Johnson, A. Karmiloff-Smith, D. Parisi and K. Plunkett. *Rethinking Innateness: A connectionist perspective on development.* MIT Press, Cambridge, MA, 1997.

[5] E. Thelen, E., G. Schoner, C. Scheier, and L.B. Smith (In press). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences.*

[6] J. Weng. The Living Machine Initiative. Department of Computer Science Technical Report CPS 96-60, Michigan State University, East Lansing, MI, Dec. 1996. A revised version appeared as a chapter: J. Weng. Learning in Computer Vision and Beyond: Development, in *Visual Communication and Image Processing*, C. W. Chen and Y. Q. Zhang (eds.), Marcel Dekker, New York, NY, 1998.

[7] N. Almassy, G. M. Edelman and O. Sporns, Behavioral constraints in the development of neural properties: A cortical model embedded in a real-world device. *Cerebral Cortex*, vol. 8, no. 4, pages 346-361, 1998.

[8] J. Weng, W. S. Hwang, Y. Zhang and C. Evans, "Developmental Robots: Theory, Method and Experimental Results," in *Proc. 2nd International Symposium on Humanoid Robots*, Tokyo, Japan, pp. 57- 64, Oct. 8- 9, 1999. An earlier result of SAIL-1 developmental program was

reported in J. Weng and W. S. Hwang, "Toward automation of learning: The state self-organization problem for a face recognizer,"in Proc. *3rd International Conference on Automatic Face- and Gesture-Recognition*, Nara, Japan, pp. 384 - 389, April 14-16, 1998.

[9] D. Roy, B. Schiele, and A. Pentland. Learning Audio-Visual Associations using Mutual Information. In Proc. *International Conference on Computer Vision, Workshop on Integrating Speech and Image Understanding.* Corfu, Greece, 1999.

[10] M. Huber and R. A. Grupen. A feedback control structure for on-line learning tasks. *Robotics and Autonomous Systems*, vol. 22, no. 3-4, pages 303-315, 1997.

[11] H. Gardner. *Multiple intelligences: The theory in practice*. Basic Books, New York, NY, 1993.

[12] N. Bayley. *Bayley Scales of Infant Development*, 2nd Edition, The Psychological Corporation, San Antonio, TX, 1993.

[13] G. Arthur, *The Arthur Adaptation of the Leiter International performance Scale*, The Psychological Service Center Press, DC, 1952.